



Contents lists available at ScienceDirect

The Journal of Systems & Software

journal homepage: www.elsevier.com/locate/jssA systematic analysis of higher education on software engineering in the Netherlands[☆]Bastiaan Heeren^a, Fabiano Dalpiaz^e, Mazyar Seraj^b, Roberto Verdecchia^{c,*}, Vadim Zaytsev^d^a Open University, The Netherlands^b Eindhoven University of Technology, The Netherlands^c University of Florence, Italy^d University of Twente, The Netherlands^e Utrecht University, The Netherlands

ARTICLE INFO

Dataset link: <https://doi.org/10.5281/zenodo.21276768>

Keywords:

Software engineering education

Systematic analysis

Higher education

ABSTRACT

Objective: Software engineering educators strive to continuously improve and refine their courses and programs. Understanding the current state of practice of software engineering higher education can empower educators to critically assess their courses, fine-tune them, and ultimately enhance their educational curricula. In this study, we provide an encompassing analysis of higher education on software engineering by considering the educational offering of The Netherlands.

Study methods: We adopt a crowdsourced analysis considering 10 Dutch universities and 207 courses. Courses are analysed via a set of key knowledge areas adapted from the SWEBOOK, which are mapped to courses by educators of their universities. The mapping process is refined via homogenisation and internal consistency improvement phases, followed by a data analysis phase.

Findings: Given its fundamental nature, Construction and Programming is the most covered knowledge area at Bachelor level. Other knowledge areas are equally covered at Bachelor and Master level (e.g., software engineering models), while more advanced ones are almost exclusively provided at Master level (e.g., Maintenance). Three clusters of tightly coupled knowledge areas emerge: (i) requirements, architecture, and design, (ii) testing, verification, and security, and (iii) process-oriented and DevOps topics. Dutch universities cover all knowledge areas uniformly, with minor deviations reflecting institutional research strengths.

Conclusions: Our results highlight correlations among key software engineering knowledge areas. We also identify underrepresented areas, such as software economics, which educators may consider including in curricula. We invite researchers to make use of our research method in their own geographical region to globally compare software engineering education programs.

1. Introduction

Software engineering (SE) education plays an important role in preparing the next generation of software professionals and researchers. Software is at the basis of all kinds of systems performing bank transactions, transportation and logistics tasks, healthcare and insurance activities. Universities want their graduates to possess solid foundational knowledge and practical skills to build and maintain software-intensive systems. Institutions of higher education thus carry a significant responsibility to provide SE courses that align with the current industry demands, technological trends and research advancements.

However, defining the contents of a SE curriculum and of SE courses entails delicate trade-offs that should account for international guidelines (Ardis et al., 2015; on Computing Curricula, 2024), standard bodies of knowledge (Washizaki et al., 2024), emerging trends (e.g., sustainability Chitchyan and Mahmoudi, 2025, artificial intelligence Hasan et al., 2024), academia-industry gaps (Akdur, 2022), and the research strengths of the academic institution (especially in post-graduate programs). The premise of our study is that *different academic institutions make different trade-offs and, therefore, we should expect visible differences*.

This paper presents the findings of a nationwide survey of SE courses offered at Dutch universities, with the goal of understanding

[☆] Editor: Gabriele Bavota.

* Corresponding author.

E-mail addresses: f.dalpiaz@uu.nl (F. Dalpiaz), m.seraj@tue.nl (M. Seraj), roberto.verdecchia@unifi.it (R. Verdecchia), v.zaytsev@utwente.nl (V. Zaytsev).

how these courses map to established bodies of knowledge and, by doing so, identifying opportunities for curriculum improvement and course sharing. This collaborative initiative stems from the VERSEN network — the Dutch National Association for Software Engineering.¹ Within VERSEN, the Working Group on Education² focuses on continuously improving SE education across the Netherlands. We study a single country because research has shown the perils of conducting cross-country comparisons in higher education (Teichler, 2014): while these would provide very valuable knowledge, there are significant contextual factors that characterise specific countries (see Section 2.1). Although a single-country study is more limited, it can be seen as a multi-case study with a relatively homogeneous sample: in our case, *research universities* in the Netherlands.

Dutch higher education is governed by *Dutch Higher Education and Research Act*, officially known as the *Wet op het hoger onderwijs en wetenschappelijk onderzoek* (WHW) (Overheid, 2025). The WHW defines the structure of university programmes and diplomas, as well as core regulations on study units and the European Credit Transfer and Accumulation System (ECTS). Under this system, Bachelor of Science (BSc, undergraduate) programmes typically consists of 180 ECTS, corresponding to three years of full-time study. These BSc programmes often culminate in a substantial project, carried out either individually or in a small group — demonstrating students' ability to synthesise SE concepts and practices. Master of Science (MSc, graduate) programmes span either one or two years (60 or 120 ECTS, respectively), depending on the university and the specific programme design. MSc programmes, besides going deeper into the content, are expected to place greater emphasis on critical thinking, research skills, and the ability to contribute new insights to the discipline (Meijers et al., 2005). Consequently, about one-quarter of an MSc student's time is dedicated to an individual final project that focuses on applying advanced topics and tackling research-oriented challenges.

Dutch universities organise the academic year into either two semesters or four quartiles. Each course confers a certain number of ECTS, and is organised around a set of competencies (generally between 5 and 7.5 ECTS), topics and learning objectives, which are explicitly listed in the course descriptions we will be analysing in the following sections. To help students navigate the landscape of courses offered within their university, BSc courses are often grouped into modules (Visscher-Voerman and Muller, 2017; University of Twente, 2023), themes, learning lines (Wijngaards et al., 2020; Cornelisse et al., 2022), etc., and MSc courses are organised into tracks or specialisations. We elaborate on it more in the next section where we introduce all the universities covered in this survey.

In this research, we classified the content of these courses against the *Software Engineering Body of Knowledge (SWEBOK) v4*, published by the IEEE Computer Society (Washizaki et al., 2024). While we acknowledge alternative frameworks such as the ACM/IEEE *Computer Science Curricula* (CS2023) (on Computing Curricula, 2024) and other discipline-specific guidelines, SWEBOK provides a widely recognised reference model specific for SE. It organises SE knowledge into distinct *knowledge areas* (KAs), offering a systematic way to map, compare, and analyse the coverage of SE topics across multiple institutions.

This article makes two contributions:

- We devise a systematic method for mapping the offering of software engineering in higher education;
- Via the application of the method to ten universities in the Netherlands, we reveal similarities and differences across knowledge areas and between universities.

Paper organisation Section 2 provides an overview of higher education in SE in the Netherlands, defining the scope of our analysis. Section 3 presents our research goal and questions, followed by our research method in Section 4. Section 5 reports our findings per knowledge area (RQ_1), while Section 6 explores correlations between these areas (RQ_2). Section 7 analyses university educational foci in SE (RQ_3). Section 8 discusses related work, and Section 9 includes a broader discussion, threats to validity, and implications for SE education. Finally, Section 10 provides a summary and directions for future research.

2. Context: Higher Education on SE in the Netherlands

We discuss higher education on SE in the Netherlands to contextualise our study. After some remarks on the general context in Section 2.1, we provide insights on BSc programmes in Section 2.2 and MSc programmes in Section 2.3, following by a definition of the scope of our analysis in Section 2.4.

2.1. General context

Dutch higher education institutions can be broadly classified into two categories: (i) **research universities** (*universiteiten*) for research-oriented higher education and (ii) **universities of applied sciences** (*hogescholen*) for higher professional education (Studiekeuze, 2025). The former category has a further sub-class, technical universities (Delft, Eindhoven, Wageningen and Twente), which have a clear focus on technical subjects and on the engineering discipline (Subramaniam et al., 2025). Unlike research universities, universities of applied sciences prioritise skills-based learning, often in collaboration with industrial partners. BSc-level graduates of research universities can freely enrol in many MSc-level programmes in their field (Overheid, 2025, Article 7.30b), while BSc-level graduates from universities of applied sciences are often required to follow a short premaster programme (six months or a year) (Overheid, 2025, Article 7.30e) to catch up on theoretical knowledge. For computer science, such a premaster typically includes several subjects in mathematics.

The landscape of SE education in the Netherlands offers many options to study seekers both from the country as well as from abroad. SE knowledge is offered by most Dutch higher education institutions, primarily within computer science programmes and departments. *Computer Science* and *Technical Computer Science* (*Informatica* and *Technische Informatica*, respectively) are established terms that attract students since the 1980s, and diplomas of such programmes are still actively sought in after the 2020s. These are complemented by programmes like *Business and IT* or *Information Science* (*Informatiekunde*), as well as variations of *Data Science* and/or *Artificial Intelligence*. There are then several smaller thematic BSc programmes like *Medical Computer Science* and *Computational Social Sciences* (both at the University of Amsterdam); similarly, at the MSc level, there are specialised programmes like *Digital Forensics* at the University of Leiden or *Computer Security* at the Vrije Universiteit Amsterdam. While this reflects the field's diversity and the willingness of higher education to be responsive to emerging technological advancements as well as the market need, in this study **we focus mostly on the (technical) computer science programmes** where SE courses play a central role.

2.2. BSc programmes

There are no specific SE programmes at the BSc level. As explained earlier, SE education is embedded in (technical) computer science programmes. The most visible distinction concerns the language of education, which can be English or Dutch. Programmes taught in Dutch are popular at institutions with an established student influx, often from within a particular region of the country, and are motivated by the perceived ease of learning in one's native language (Soosai Raj et al., 2018) as well as the regional job market needs. Examples are the

¹ <https://www.versen.nl>

² <https://www.versen.nl/contents/works/education>

computer science programmes of Open University, Utrecht University, University of Amsterdam, and Leiden University, having cohort sizes of 148 to 221 students in 2024 (Keuzegids, 2024b). Programmes in English are motivated by attractiveness for international applicants, but also appealing for Dutch students aiming at international careers. Universities like Radboud University in Nijmegen, University of Groningen, Vrije Universiteit Amsterdam, Eindhoven University of Technology, and University of Twente teach and run their programmes in English. These programmes often see a broader influx, such as Vrije Universiteit Amsterdam, which enrolled 542 students this year (Keuzegids, 2024b). The smallest programme among these is the one offered in Groningen, with 121 students, the other programmes are at 204–289 students (Keuzegids, 2024b). The Technical University of Delft recently took the decision to take a hybrid approach, offering parallel influx tracks in Dutch and English, attracting 474 students overall (Keuzegids, 2024b). The decision to teach in English often hinges on strategic factors such as long-term growth, and even though the cultural shift associated with English-language instruction and the multiculturalism that brings its own challenges, it does simplify hiring of teaching personnel.

Some BSc programmes limit student intake via the so-called “*Numerus Fixus*”, an admission cap based on an entrance exam result, regulated by a separate law, called *Regeling aanmelding en toelating hoger onderwijs* (Regulations for Registration and Admission to Higher Education) (Overheid, 2024). This measure places a legal limit on enrolment numbers, and students are only invited to join the programme in the order of their ranking based on their performance on the entrance exam. Institutions such as University of Groningen, Vrije Universiteit Amsterdam, Eindhoven University of Technology, and Delft University of Technology have adopted this approach to control intake and maintain quality standards in the face of rising demand. At the BSc level, this is the only admission and selection procedure which allows a university to guarantee a cap on their student influx.

2.3. MSc programmes

MSc programmes across the Netherlands range from general computer science curricula to highly specialised offerings. General computer science programmes, offered by most universities, typically provide tracks or specialisations³ that are more closely aligned with SE, such as *Software Science* at the University of Groningen, *Software Technology* at the University of Twente and *Software and Analytics* at Eindhoven University of Technology. These tracks or specialisations enable students to focus on core SE competencies within a broader computer science framework, allowing for the integration of SE principles with foundational computer science. Yet, specialised programmes such as the University of Amsterdam’s or Open University’s MSc in *Software Engineering*, aim to provide a deeper, more targeted curriculum that aligns with industry trends in automation, agile practices, and advanced software testing methods.

Several institutions offer MSc degrees in *Business and IT* or related somewhat interdisciplinary programmes, bridging technical and business knowledge. These programmes are officially classified and evaluated in a separate educational category (*Information Science*).

Enrolment sizes across Dutch MSc programmes vary significantly, reflecting differences in institutional focus, student demographics, and intake flexibility. Delft University of Technology has the largest related MSc programme, with 258 new enrolments in 2024, according to Keuzegids data (Keuzegids, 2024a), while the joint MSc programme between the University of Amsterdam (UvA) and Vrije Universiteit Amsterdam (VU) is the second-largest, inviting 151 students on the same year (Keuzegids, 2024a). Leiden University, which does not have

a dedicated SE track, reported 126 new students in its general computer science program, making it the third-largest. Mid-sized programmes, with enrolments ranging between 69 and 112 students, include those at Twente, Eindhoven, Utrecht, Nijmegen, and UvA’s specialised *Software Engineering* MSc (Keuzegids, 2024a).

Note that enrolment figures are challenging to define accurately especially in MSc programmes, as students often have multiple entry points and pathways. For example, in Utrecht and Twente, new students can start in either September or February. Additionally, unlike BSc programmes, where students progress together as a cohort, MSc students frequently graduate on individual timelines which reflect their personal study experience.

At the MSc level, SE education in the Netherlands follows a remarkably international orientation, with English serving as the dominant language for study materials, which is also in general typical for computer science education. This occurs largely due to the recognition that fluency in English, both in technical vocabulary and in day-to-day communication, is an essential skill for future experts to stay competitive on the global SE market. The noticeable exceptions are the Open University’s two distinct MSc programmes in *Computer Science* as well as in *Software Engineering*, which are taught in Dutch, catering specifically to a domestic audience (as well as covering Dutch-speaking regions of neighbouring Belgium) and offering some flexibility for Dutch-speaking professionals.

2.4. Scope of our analysis

Our research focuses on SE courses offered within computer science and SE programmes offered by the following research universities in the Netherlands, listed in alphabetical order, with names of cities that host them, in parenthesis, and abbreviations that we will use from now on to save space:

- Delft University of Technology (Delft) — TUD
- Eindhoven University of Technology (Eindhoven) — TU/e
- Open University (Heerlen) — OU
- Radboud University (Nijmegen) — RU
- Rijksuniversiteit Groningen (Groningen) — RUG
- University of Amsterdam (Amsterdam) — UvA
- University of Leiden (Leiden) — UL
- University of Twente (Enschede) — UT
- Utrecht University (Utrecht) — UU
- Vrije Universiteit Amsterdam (Amsterdam) — VU

We therefore exclude Maastricht University, which did not offer a computer science programme at the time this analysis was performed; Erasmus University Rotterdam, which profiles itself in economics and medical sciences and thus does not provide SE education; Tilburg University, which similarly focuses on humanities and law; and Wageningen University, which focuses on agricultural and environmental research and education. We also exclude all universities of applied science (at least 48) not only to make the work feasible, but also to avoid the discussion on the level of the learning objectives, since they have a more applied and market-driven nature.

3. Research goal and questions

In order to define our goal and research questions, we follow the Goal-Question-Metric approach first suggested by Basili et al. (1994). Our overarching research goal is defined as follows:

Analyse the SE educational landscape
 For the purpose of understanding the trends
 With respect to course content
 From the viewpoint of SE educators
 In the context of the Netherlands.

³ Tracks are ministry-regulated, and specialisations can have different definitions per university (van den Broek and van Casteren, 2018).

As outlined in our goal, the objective of this research is to gain a systematic understanding of the current higher education on SE in the Netherlands. As further detailed in Section 4.2, such goal is achieved via a systematic analysis of course content crowd-sourced across educators of 10 Dutch universities providing SE academic instruction.

Following the GQM approach, we derive the following research questions (RQ), that we need to answer in order to achieve our research goal.

RQ₁: What are the topics studied in SE Courses?

We aim to gain insights into the content of the SE courses provided by the considered Dutch universities. As further specified in Section 4.2, to answer such RQ, we focus on the analysis of key *Knowledge Areas* (KAs), initially taken from the SWEBOK, and further refined via an iterative process. The data collected for RQ₁ lays the groundwork to answer our subsequent RQs.

RQ₂: What are the (co-)occurrences among studied topics and the educational context?

We investigate the potential inter-dependencies between SE topics, not only by considering inter-KA occurrences, but also by taking into account the frequency with which the topics appear among graduate and undergraduate courses. In addition, for RQ₂, we also study the role that SE-specific programmes and tracks play in the presence of KA in courses.

RQ₃: Do universities differ in terms of education foci?

We look for differences in terms of SE educational content across the considered universities, which we would expect because of the different trade-offs universities make when defining their programs and courses. By answering RQ₃, we can assess the extent to which educational programmes providing SE education display heterogeneity and similarity across their formative offer.

4. Research method and process

We first discuss the broader research methodology aspects of our study in Section 4.1, and then detail the research process we followed in Section 4.2.

4.1. Methodological embedding

Our research method is a combination of: (i) *cross-case analysis* (Bor-man et al., 2012), as we are comparing multiple universities; and (ii) *deductive qualitative analysis* (Gilgun, 2019), as we analyse collected data in a predominantly qualitative manner according to an existing theory/framework: the knowledge areas in the SWEBOK. Although these techniques are often used within grounded theory (Glaser et al., 1998), we cannot argue we are conducting that research method because we do not construct a theory as part of our analysis.

Our study can be seen under the lens of *comparative higher education research* (Teichler, 2014). The field, as observed by Teichler (2014), is characterised by an *inherent tension* between obtaining a rich contextual understanding (best addressed in a single institution or country) and performing a comparison at the international level. Teichler explains how national studies are often seen as “over-descriptive”, while international studies tend to be shallow because of limited contextual knowledge, they involve a single researcher per country, they suffer from a coincidental choice of countries, often a single foreign (conveniently selected) country. Indeed, Kosmützky and Krücken (2014) found that 67% of the international studies compare only two countries. The debate on how to best tackle the tension between context and comparability is still open (Kosmützky et al., 2020).

This study is a first step towards *theory building*: our analysis offers a replicable process (see Section 4.2) and the key results that we report in this paper are highlighted as a set of *findings*. These findings derive from a national context that involves every research university in the Netherlands with a focus on SE education. We involve multiple researchers from one country, based on their profound knowledge of the regional context (Section 2), and we crowdsource data collection to the universities themselves. We acknowledge that *theory testing* is an important future direction: to what extent are our findings specific to the context of the Netherlands, to the time snapshot of our analysis, and to other factors? We encourage researchers from other countries to replicate the study either at the national or international level; in the latter case, particular attention must be placed on achieving rigour in the selection of countries and institutions (Kosmützky, 2016).

4.2. Research process

The research process we followed, depicted in Fig. 1, is composed by the following steps:

1. *Define initial KAs and keywords.* Two of the five authors of this paper have initiated the process by identifying a reference framework for categorising the courses according to knowledge areas of SWEBOK. At the time of our data analysis, SWEBOK was in a phase of transition from version 3 to version 4. As such, these authors considered the structure and a draft of the SWEBOK V4, which provided a more up-to-date view of the discipline⁴ including 18 knowledge areas (KAs) than SWEBOK v3 which dates back to 2014. Out of these 18, the authors selected the 15 key KAs: Requirements, Architecture, Design, Construction, Testing, SE Operations, Maintenance, Configuration Management, SE management, SE process, Models and Methods, Quality, Security, Professional Practice, and SE economics. We excluded the KAs that represent foundational knowledge: Computing Foundations, Mathematical Foundations, and Engineering Foundations. For each of the 15 KAs, the authors have read the draft chapter of SWEBOK V4 and have identified a set of keywords that are likely to characterise key topics within the KA. For example, for the KA Requirements, we identified Requirements, ‘Elicitation’, ‘Specification’, ‘Analysis’, ‘System objectives’, ‘Use cases’, ‘Non-functional requirements’, etc.
2. *Annotate initial set of courses.* Three authors (the two involved in step (1) and an additional one) collected and annotated — via a spreadsheet — collections of SE-specific courses from four universities: OU, TU/e, UU, and VU. These were gathered either directly by the authors being an employee of a university, or through an examination of the course catalogues with the help of a colleague working at that institution. The annotation process consisted of: (a) reading thoroughly the course description; (a) ticking the KAs that the course contributes to, based on the presence of the keywords for that KA or related terms, and (c) taking notes about the course or the annotation process, when necessary. We did not employ multiple annotators for these 85 courses; however, the three authors checked each others’ annotations and had multiple rounds in which they discussed discrepancies, leading to the next step.
3. *Refine KAs and keywords.* Based on the experience gained from the annotation process and from the discussion meetings, the three authors involved in step (2) worked at a revised set of KAs and keywords. A few major changes were applied:

⁴ For examples, SWEBOK V4 decouples software architecture from software design, a distinction that has been argued for since the 1990s (Perry and Wolf, 1992).




Fig. 1. Research process overview.

- Given the high number of courses related to programming, we have extended the Construction KA to Construction & Programming.
- Since the methodological component of the KA Models & Methods overlaps with SE management and SE process, we restricted the KA to modelling and model-driven engineering, and we use the term SE models.
- We removed two KAs: Software quality, as it can be argued that most KAs contribute to software quality; and Professional practice, as this was found too seldom in the examined courses.
- We introduced two additional KAs that represent an important aspect of SE education in the Netherlands and that are not listed as KAs by the SWEBOK: Verification and Programming language design (briefly: PL Design).
- We introduced an additional field that denotes the type of course: (a) Programming denotes a course teaching a programming paradigm or language; (b) PL Design groups the courses that focus on the design of programming languages; (c) Seminar indicates research courses; (d) SE101 refers to introductory courses on SE; (e) Project indicates courses where the learning is mainly achieved by conducting a project (Gary, 2015) for a real or simulated client.
- We removed six courses (from the 85 we analysed) that were deemed irrelevant.

Table 1 provides an overview of the final annotation scheme; the top part shows the KAs, while the bottom part focuses on the clusters of courses.

4. *Extend the set of annotated courses.* The revised annotation scheme from step (3) was then used in a second annotation round. The three authors involved in step (2) have re-analysed the courses from the four universities, and at the same time, selected researchers/educators from six additional universities (RU, RUG, TUD, UvA, UL, UT) were invited and asked to (a) provide a list of SE courses taught at their institution, and (b) annotate those using the annotation scheme. This led to an additional set of 207 courses with annotations. These additional colleagues were selected based on convenience sampling: we invited contacts we had met at events on SE education and with expertise beyond teaching individual courses.

5. *Homogenisation of course annotations.* The new set of 207 courses was analysed by the same three authors in order to identify discrepancies from the annotation that was performed on the 79 courses (which were left over from 85 when six were removed). We observed that many courses were prerequisites for SE courses but could not be mapped to the KAs we included; this led to removing 103 courses, thereby retaining 128 courses for the additional six universities. Summing up to the annotated courses from the first four universities, this results in 207 courses that are included in this analysis. During this process, the three authors also checked the ticked KAs against the course descriptions, and made several adjustments whenever they saw a discrepancy from the annotation scheme. We have found, indeed, that several KAs had been ticked even when that KA was not a primary focus for that course. In total, this resulted in 167 removed ticks and 39 added ones. Given the 319 agreements to tick and the 2580 agreements to not tick the KA, this results in a Cohen's kappa of 0.718, which represents substantial agreement. A detailed analysis of the inter-rater agreement per category can be found in the online appendix.⁵ We sent the list of retained courses and our revised annotations to the colleagues who provided the initial lists, giving them the opportunity to challenge our decisions; we have not received objections, only a couple requests for clarifications.
6. *Improve internal consistency.* In order to improve the internal consistency of the tagging, the three authors subdivided the KAs and they checked all 207 courses, with the aim of determining the consistency of the annotations for each KA. The goal was that of ensuring that a single annotator would make a 'vertical' analysis of all courses against one individual KA, looking for omitted or erroneous ticks. This led only to three adjustments: one tick was added, two were removed. We also identified *signature courses*; we labelled a course as *signature for a KA* if the course revolves almost exclusively around that KA (e.g., a 'Requirements Engineering' course would be a signature course for the KA Requirements).
7. *Analyse the data.* The resulting spreadsheet was used at the basis of the follow-up analysis, which revolved around the research questions described in Section 3.

⁵ Our online appendix includes our coded data as well as other scripts used for our analysis: <https://doi.org/10.5281/zenodo.21276768>.

Table 1

Overview of the final annotation scheme. The upper part shows KAs, the lower part focuses on course clusters.

Knowledge area	Keywords/topics
Requirements	requirements, elicitation, specification, analysis, system objectives, goals, user stories, use cases, non-functional requirements
Architecture	software architecture, ADL, 4+1 model, concerns, architectural styles, software product lines, architectural definition, architectural patterns, quality attributes, architectural tactics
Design	UML, class diagrams, object-oriented design, software design, design patterns, refactoring, unified process, domain modelling
Construction & Programming	programming (imperative, object-oriented, functional, web, game, model-driven), C#, Java, Haskell, Scala, Python, HTML, PHP, Javascript, app development, business rules, code documentation, IDE
Testing	testing (unit, integration, system, acceptance), property-based testing, record-and-replay, mutation testing, partition testing, white box, black box, test cases, test set, coverage criteria, testing model, software quality, scriptless, model-based
SE Operations	continuous delivery, continuous integration, deployment, integration, packaging, monitoring, release engineering
Maintenance	software evolution, maintenance, change, product quality, software quality management, impact analysis, technical debt, software re-engineering
Configuration management	version control, git, svn, configuration, dependencies, source code management system
SE management	SDLC, project planning, cost estimation, effort estimation, requirements change, schedule, SE measurement
SE process	(product) life cycle, methods, waterfall, iterative, agile, spiral, V-model, Scrum, software process assessment and improvement, minimum viable product, proof of concept
SE models	domain modelling, UML, variability modelling, OCL, business modelling, metamodelling, process modelling, BPMN, model-driven development
Software Security	security vulnerabilities, security by design, sandboxing, common vulnerabilities, security risk analysis, ethical aspects of security, penetration testing, security patterns
SE economics	finance, accounting, controlling, cash flow, valuation, product vs. project vs. system, costing, estimation, prioritisation
Verification	verification, LTL, CTL, model checking, symbolic execution, theorem proving, formal specification, static analysis, program analysis
Progr. Language design	program syntax, semantics, compilers, types, type systems
Course cluster	Description
SE 101	A basic first course on software engineering
Project	A project-based course where the students learn, either solo or in teams, how to apply SE in a practical setting. Often a capstone project.
Seminar	A research-oriented course on advanced SE topics.
Programming	A course that teaches a programming paradigm or language.
PL Design	A course on the design of programming languages.

Table 2

Top twenty common terms for the KAs Requirements, Architecture, Design, Testing, and Maintenance with relative frequency compared to the most frequently occurring term in that KA.

Requirements	Freq.	Architecture	Freq.	Design	Freq.	Testing	Freq.	Maintenance	Freq.
requirement	1.000	architecture	1.000	design	1.000	testing	1.000	software	1.000
technique	0.708	software architecture	0.791	software	0.958	test	0.544	software evolution	0.548
specification	0.583	software	0.674	model	0.583	software	0.533	software system	0.452
software	0.542	system	0.535	system	0.552	system	0.478	quality	0.355
system	0.458	design	0.442	architecture	0.396	technique	0.444	metric	0.290
user	0.375	requirement	0.279	analysis	0.344	model	0.389	technique	0.258
need	0.375	concern	0.209	engineering	0.333	software testing	0.233	code	0.258
elicitation	0.333	stakeholder	0.209	requirement	0.313	code	0.189	tool	0.226
requirement engineering	0.333	quality	0.209	use	0.292	quality	0.189	engineering	0.194
diagram	0.292	project	0.186	tool	0.281	risk	0.178	maintenance	0.194
process	0.250	architect	0.186	problem	0.260	program	0.167	requirement	0.194
practice	0.250	structure	0.163	application	0.260	use	0.156	language	0.194
development	0.208	model	0.163	software system	0.260	design	0.144	method tool	0.194
use	0.208	component	0.140	design pattern	0.250	process	0.133	evolution	0.161
formal	0.208	practical	0.140	method	0.229	tool	0.122	change	0.161
method	0.167	style	0.140	principle	0.229	specification	0.111	development	0.161
informal	0.167	decision	0.140	project	0.219	analysis	0.111	program comprehension	0.161
project	0.167	concept	0.140	software development	0.219	state	0.100	evolution discuss	0.161
activity	0.125	development	0.140	programming	0.198	method	0.100	step	0.129
stakeholder	0.125	analysis	0.140	quality	0.198	modelbased testing	0.100	analysis	0.129

5. Findings per knowledge area (RQ₁)

We present findings per each knowledge area we analysed, discussing the number of courses per KA, their distribution across BSc and MSc programs and tracks, course type, and an analysis of the course topics through an analysis of the most frequently occurring topics. For this last activity, we used the wordcloud package⁶ to generate the top 20 terms in the KA, also including two-word compounds, after removing stopwords, and applying lemmatisation. For the sake of brevity, we present A summary of the recurrence of the KAs across courses is shown in Fig. 2.

5.1. Requirements

Looking at the coverage of the Requirements KA, we see that nine of the ten universities considered cover this knowledge area in their courses, offering a total of 20 courses. Twelve of these courses are offered to BSc students. Eight courses are taught in SE-specific programs or tracks. Three of them are at UvA as part of their MSc in *Software Engineering*. All 20 courses are regular ones, with no specific course type such as ‘Seminar’ or the like.

Five universities (RU, UL, UU, UvA, and VU) have signature courses that include the term *Requirements Engineering* (RE) in the course name. Among the other universities, TU/e has a highly related course called *Software Specification*. Out of these six courses, four are taught to

⁶ <https://pypi.org/project/wordcloud/>




Fig. 2. Recurrence of topics (KA) treated in the courses.

undergraduate students, and two (UU, UvA) are offered to graduate students.

The first column of Table 2 presents the most frequent terms and shows that, in addition to the obvious term ‘requirement’, highly prevalent topics refer to the fact that ‘techniques’ are being taught (as opposed, e.g., to algorithms or theories), the ‘specification’ phase, the focus on ‘software’ as well as on ‘systems’. The words ‘user’ and ‘need’ highlight the importance of considering the perspective of the actual users of the system, and the fact that these have needs, rather than ready-made requirements. Other interesting terms are ‘elicitation’ and ‘diagram’, classical phases and techniques that are taught in these courses.

5.2. Architecture

There are 22 courses offered by the ten examined universities; just like for requirements, all universities offer at least one course that covers this KA. However, only four courses are at the BSc level, indicating that software architecture is deemed as an advanced topic in the context of SE education. 15 courses are part of an SE-specific track or program.

There are six signature courses that focus solely on software architecture (OU, RUG $\times 2$, TUD, UT, UU). They are all at the MSc level, also indicating the advanced nature of the topic in the analysed curricula. Interestingly, RUG offers two courses on software architecture; one of which, denoted as *advanced*, has a clear focus on preparing professionals, also thanks to the emphasis on a project. Two courses are offered as research seminars: *Business Information Systems* (UL) and *Fundamentals of Adaptive Software* (VU).

Some of the analysed courses take a perspective that goes beyond pure software, and they enter the field of information systems. While software architecture is a core topic in SE, architectural design is also a prominent activity for other design artefacts. This is visible, for example, in the following courses: *Seminar Business Information Systems* (UL), *ICT Architectures* (UL), *Architectures of Information Systems* (UT), and *Digital Architecture* (VU).

The second column of Table 2 highlights that, in addition to the obvious keywords ‘architecture’, ‘software architecture’, ‘software’, ‘system’, and the aforementioned ‘design’, the following most prominent topic is that of ‘requirement’. This witnesses the relationship between the two knowledge areas (Nuseibeh, 2001). Other frequent terms are ‘concern’, ‘stakeholder’, ‘quality’, ‘structure’. These clearly denote key terms in the field such as stakeholders’ concerns, architecture as a structure, and quality aspects.

Finding 1: While equally covered overall, the Requirements KA is covered both at the BSc and MSc level, while the Architecture KA is mainly included in MSc level courses, indicating that this KA is probably seen as a more advanced topic by curriculum designers. This is in line with observations that teaching this KA is difficult (Galster and Angelov, 2016).

5.3. Design

There is a total of 24 courses that are offered by nine of the ten universities in our list. Half of these are at the BSc level, half at the MSc level. Seven universities have at least one BSc course, seven have at least one MSc course. Only four universities (OU, RUG, UT, VU) have dedicated courses on this topic, which links early phases like requirements engineering and software architecture with programming, at both levels. Interestingly, three courses are in the cluster ‘Programming’, indicating that design and programming are sometimes taught together.

Defining signature courses for this category is not simple, as Design is often taught in conjunction with other courses. Six highly specific courses can be identified, out of which three have a clear focus on this category: *Software Design* (TU/e and VU) and *Problem Analysis and Software Design* (RUG). Other three courses focus on advanced design aspects: *Design for Change* (OU), *System Design Engineering* (TU/e), and *Service Oriented Design* (VU). Half of these courses are part of SE programs or tracks.

When we explore the most frequent terms in the descriptions of the courses in this category (Table 2, third column), we see that the term ‘software’ stands out much more prominently than for the previously analysed KAs Requirements and Architecture. Other major keywords are obviously ‘design’ itself, ‘model’ (linking to the SE models KA), ‘system’ (indicating a broader perspective than software only), ‘architecture’, ‘analysis’, ‘engineering’, and ‘requirement’.

5.4. Construction & programming

This KA is the most frequent: a total of 75 courses across all ten analysed universities cover this KA. The undergraduate courses are by far more popular: 62 of the 75 courses are at the BSc level, only 13 are at the MSc level. This is in line with the key idea that learning programming is a core competence of any computer science student and, as such, it has to be covered at the BSc level. This is further confirmed by the fact that only 12 courses are part of a SE-specific track or program. Within the 75 courses, 55 are also in the cluster ‘programming’, denoting a course with the main goal of learning a programming paradigm or language.

By analysing the 13 MSc level courses, we observe that these cover advanced programming paradigms, including parallel programming (RUG and VU), web and cloud computing (RUG), cryptography and security (TUD, UvA, UL, UT), and functional programming (UU, RU, UvA).

Finding 2: At the basis of SE education lies Construction & Programming, which appears in most BSc SE degrees (63/75 occurrences) in courses that focus exclusively on programming (55/75). As a foundational aspect, programming syntax is the first stepping stone towards an educational path into SE knowledge.

Given the foundational role of Construction & Programming for computer science and SE, we propose that further studies could be conducted with a sole focus on analysing how programming is being taught across universities, in line with previous studies on the choice of the programming language for introductory courses (Murphy et al., 2016; Simon et al., 2018).

5.5. Testing

In total, 19 out of the 207 total courses considered present a focus on the Testing KA. In terms of universities, Testing courses are offered in the vast majority of the universities included in this study (nine out of ten). Similar to the Architecture KA (Section 5.2), most courses with an emphasis on software testing are provided at the MSc level (12 out of 19), while only a third at BSc level. This could be attributed to the impossibility to dig deep into testing concepts while conducting a general purpose BSc programming course, leading to the higher presence of the KA in more advanced courses, which necessarily need to be provided at the MSc level.

Many courses (11/19) are part of a SE-specific track or program, indicating that this KA is highly specific to SE. Most courses are regular ones, with three exceptions: one seminar (*System and Software Security* at UL), one in the programming cluster (*Cryptographic Engineering* at UvA), and one project (*Project System Testing* at VU). The latter two highlight that some educators emphasise how the value of testing is best understood by practising skills.

Finding 3: Like Architecture, Testing is also taught mostly at the MSc level, and mainly in SE-specific tracks and programs, signalling a highly specific KA for SE. This aligns with a previous study (Tramontana et al., 2024) that shows how software testing courses in Belgium, Italy, Portugal and Spain are mostly taught at the MSc level, with a few exceptions in the final BSc years.

By considering signature courses, we notice these to be the majority (12 out of 19). Most are provided at the MSc level (9 out of 12), further corroborating that courses putting high emphasis on testing require a more advanced understanding of SE practices, and hence are considered mostly for MSc academic curricula.

The most frequent keywords in the course descriptions (fourth column in Table 2) are the obvious ‘testing’, ‘test’, ‘software’, and ‘system’. Other frequent keywords include ‘technique’, as testing is often taught as a set of techniques that can be used; ‘model’, probably referring to model-based testing; ‘code’ and ‘program’, indicating against which artefact tests are run; and ‘quality’ and ‘risk’, as testing is proposed as a component of quality assurance (Naik and Tripathy, 2011) that can reduce the risk of failure.

5.6. SE operations

The SE Operations KA is covered in 15 out of 207 total courses, and is considered to different extents in the educational programmes of all ten universities considered for this study. The university that provides the higher number of courses with an SE Operations component is UL with two BSc and two MSc courses considering the topic. In terms of BSc and MSc courses, we note that the KA is more recurrent at MSc level (9 out of 15 total courses with a component focusing on SE Operations), while being also quite recurrent at BSc level (6 out of 15). Circa half of the courses (7) are part of SE tracks or programmes. Notably, SE Operations are taught in a variety of methods, including an SE101 course, a project, and two seminars.

The courses in the SE Operations KA are seldom considered as signature courses (2 out of 15), exceptions being *Software Containerisation* at VU and *DevOps and Cloud-Based Systems* at UvA. Given the low number of signature courses, we do not create a word cloud and we do not analyse the recurring keywords, as they would be influenced too heavily by a single course.

5.7. Maintenance

Out of the 207 courses considered in this study 19 focus on the Maintenance KA. In terms of universities, courses with a maintenance components are provided in seven out of ten universities, with TUD

resulting the university with more courses focusing on the topic (4 Maintenance courses). Almost all courses are provided at the MSc level, with only one course, namely *Software Engineering Methods* at TUD, considering the Maintenance KA at the BSc level.

The signature courses in this knowledge area are tightly coupled with the software evolution topic, and are provided at five different universities: RUG, TU/e, UL, UT, and UvA.

Given the specialised topic, which could be considered as rather specific to SE, it comes at no surprise that the majority of Maintenance courses are provided within SE-specific programmes and tracks (15/19).

A minority of courses utilises seminars as learning objectives evaluation (6/19), while only one course, namely the already mentioned *Software Engineering Methods* at TUD, has a maintenance component as part of an introductory course on SE.

When we examine the most common words in the signature courses (Table 2, fifth column), we observe that ‘software’, ‘software evolution’ and ‘software system’ are by far the most significant words, followed by ‘quality’ (maintenance is advocated as a key determinant of software quality), and ‘metric’ (used to measure quality). The following keywords refer to the use of ‘techniques’, the relationship to ‘code’, and the existence of ‘tools’ to perform maintenance.

5.8. Configuration management

From our analysis, Configuration Management results overall an infrequent KA, as it is covered in only 9 out of the 207 courses. It is considered in six universities, with RUG being the only one reporting two different courses focusing on such KA.

A balance exists between BSc (4/9) and MSc courses (5/9). Given the specificity of this KA, the inexistence of signature courses is not surprising. Given the overlap with other KAs (see Section 6), we do not analyse the most frequent terms.

One of the courses is a seminar, two courses are part of the ‘programming’ cluster, and over the half of the courses (5/9) are taught in SE-specific tracks or programs.

5.9. SE management

Unlike the previous KA, SE management emerges as a rather popular KA, covered in 24 courses. All but one university consider such topic, with RU, RUG, and UL having four courses each. Only a fraction of BSc courses have this KA (6/24). From an inspection of the BSc courses, they all seem to provide a broad overview of SE practices, such as *Software Engineering Methods* at TUD and *Software Process* at UvA.

Five signature courses on SE management could be identified: *Software Asset Management* at VU, *Systems Development Management* at RU, *Software Project Management* at TU/e, *Software Management* at UT, and *Project Management* at UL (the only signature course at the BSc level).

Eleven courses are provided in the context of SE-specific tracks and programs, showcasing the considerable specialisation on SE knowledge the KA implies.

While the majority of the courses considering the SE management KA relies on written exams, we also note a slight heterogeneity of learning objectives evaluation methods, with some courses utilising seminars (4/23), projects (2/23), or being part of SE introductory courses (1/23).

When we analyse the most frequent terms (Table 3, first column), in addition to the obvious terms ‘software’ and ‘management’, the keyword ‘process’ stands out to indicate the link with managing projects. These are followed by ‘development’ (this is the main ‘process’ being managed), and ‘model’ (referring to maturity and measurement models).

Table 3

Top twenty common terms for the KAs SE management, SE models, Software Security, SE economics, and Verification with relative frequency compared to the most frequently occurring term in that KA.

SE management	Freq.	SE models	Freq.	Software Security	Freq.	SE economics	Freq.	Verification	Freq.
software	1.000	system	1.000	security	1.000	software	1.000	program	1.000
project	0.800	model	0.786	software	0.520	development	0.431	system	0.500
management	0.500	software	0.595	system	0.280	product	0.379	technique	0.424
development	0.400	engineering	0.381	analysis	0.260	management	0.362	software	0.390
model	0.367	development	0.333	vulnerability	0.240	software product	0.310	security	0.339
process	0.300	tool	0.333	secure	0.240	software ecosystem	0.310	model checking	0.314
quality	0.200	design	0.333	computer	0.220	business	0.276	use	0.288
measurement	0.200	information	0.310	software development	0.200	model	0.224	testing	0.280
issue	0.167	model-driven	0.286	block	0.180	process	0.207	specification	0.280
team	0.167	language	0.286	programming language	0.160	technique	0.190	type	0.280
risk	0.167	transformation	0.238	level	0.140	software vendor	0.190	tool	0.271
agile	0.133	cps	0.214	technique	0.140	project	0.172	model	0.271
system	0.133	business rule	0.214	principle	0.120	technology	0.172	programming language	0.263
manager	0.133	modelling	0.190	static	0.120	market	0.155	verification	0.220
product	0.133	pattern	0.190	code	0.120	sustainability	0.155	property	0.212
metric	0.133	complex	0.190	problem	0.120	business model	0.155	formal	0.212
standard	0.133	problem	0.190	application	0.120	method	0.138	proof	0.203
insight	0.133	cyberphysical system	0.190	attack	0.120	topic	0.138	correct	0.195
cost	0.133	application	0.167	issue	0.100	measurement	0.121	logic	0.195
effort estimation	0.133	specification	0.143	type	0.100	industry	0.121	prove	0.195

5.10. SE process

The KA SE process is quite popular among SE courses, with 25 out of 207 courses focusing on such KA, and all universities providing at least two courses on it.

The KA is most treated as part of MSc curricula (17/25), while only a single MSc course, namely the *Software Process* provided at UvA, is a signature course on the topic.

Slightly more than half of the times (15/25), the KA is considered as part of a SE-specific track or program, showing how the topic is important also for students who do not explicitly study SE.

Finding 4: SE management and SE process are well-covered KAs (24 and 25 courses, respectively): SE management is offered in all but one university, SE process by all universities. This indicates that solid attention is paid to explaining the non-technical aspects of SE. We observe that this is supported by the fact that most popular SE textbooks (Pressman and Maxim, 2019; Sommerville, 2015) include ample coverage of these topics.

5.11. SE models

As explained in Section 3, the KA SE models is interpreted here in terms of the use of modelling languages and notations within SE courses.

This is a popular KA, with 28 identified courses, evenly split between BSc and MSc. Nine of the ten analysed universities offer at least one course in this category; all of them at the MSc level, seven of them also through BSc courses.

We could identify five signature courses; in addition to two having obvious names: *Model-Driven Development* (BSc, OU); *Model-Driven Engineering* (MSc, UT), we could also find one BSc course on *Software Design* (VU), as well as MSc courses called *System Design Engineering* (TU/e) and *Model-Based Design of Cyber-Physical Systems* (UvA). All of these signature courses also tick the KA Design, showing the tight interplay between the KAs.

Among these signature courses, the TU/e course that emphasises modelling in the context of systems engineering and the UvA course on cyber-physical systems, both emphasising how models are an important abstraction not only for software, but for systems in general.

Finding 5: SE models seems to be a highly represented KA compared to others. This could be explained by the strong focus on model-based SE in the Netherlands, evidenced, e.g., by the involvement of Dutch scholars in the definition of a body of knowledge for model-based SE (Burgueño et al., 2019).

This observation is confirmed by the most occurring terms in the signature courses (Table 3, second column), where ‘system’ appears as the most frequent term. The second most frequent term is obviously ‘model’, followed by ‘software’ and ‘engineering’, ‘development’, ‘tool’, all referring to aspects of modelling paradigm (model-driven engineering (MDE) and development (MDD)). Then, keywords like ‘design’ and ‘information’ follow, denoting the link with design activities and the necessity of modelling information. The next three terms ‘model-driven’, ‘language’, and ‘transformation’ are clearly referring to how languages and transformations are at the heart of MDD and MDE.

5.12. Software security

The KA Software Security includes a total of 16 courses that are taught by nine of the ten covered universities. The topic seems to be more popular at the MSc level (11 courses) than at the BSc level (5 courses). Only two universities (UL and UU) offer courses in this KA at both levels.

There are five courses that can be considered signature; three have the same name as the KA and are taught at OU, RU, and UT; the *Security* course at UL which has a clear focus on system and software design; and the *System and Software Security* course at UL, which is an advanced version of the former. Out of the various courses, only two are offered as seminars: *Security Verification* (UT) and *System and Software Security* (UL).

When analysing the most common keywords (Table 3, third column), we observe how — after the words that identify the KA itself and the typical word ‘system’ — we find the generic term ‘analysis’, the keyword ‘vulnerability’ occurs, one of the key terms in the field. The following keywords based on frequency highlight basic facets of the KA: ‘secure’, ‘programming’, ‘software development’, and ‘programming language’.

5.13. SE economics

The KA SE economics is not popular; only nine courses have a focus on this KA. These are taught across seven universities, although

one (UU) has a prevalence with three courses. Two courses are taught at the BSc level, the other seven are at the MSc level. Furthermore, two of the MSc courses are seminars, indicating that this is a topic that is still less mature when it comes to its teaching in higher education (although this is a well established topic in industry). Six of the nice courses are part of a SE-specific track (see [Boehm and Sullivan, 2000](#); [Burgueño et al., 2019](#); [Galster and Angelov, 2016](#); [Pressman and Maxim, 2019](#); [Sommerville, 2015](#); [Tramontana et al., 2024](#)).

Finding 6: SE economics is the least popular KA in our analysis. This aligns with the early observations by Barry Boehm (Boehm and Sullivan, 2000) and the limited embedding in standard curricula; the 2014 ACM SE curriculum recommendations indicate a total of *only* eight hours for this KA (Ardiis et al., 2015).

Analysing the most frequent keywords in the course descriptions (Table 3, fourth column), we see ‘software’ appearing as the top, unsurprisingly. The following keywords, with the exception of ‘development’, are more interesting: they refer to ‘product’, ‘management’, ‘software product’, ‘software ecosystem’, and ‘business’. These have clear links to research domains that are highly linked to industrial needs, including software production (Xu and Brinkkemper, 2007), software business (Cusumano, 2008), and software ecosystems (Manikas and Hansen, 2013).

5.14. Verification

The KA Verification, which we added to those from the SWEBOK because of its prominence in the SE landscape in the Netherlands, totals 28 courses taught by all the ten universities under analysis. The courses regarding verification are mostly taught in MSc programmes (18 courses), with ten courses at the BSc level. It is worth noting that 8/10 universities include Verification courses at the BSc level, with the exceptions of OU and UvA. All ten universities offer at least one course at the MSc level. Almost half of the courses (14) are part of a MSc or a track that is SE specific. Three courses are given as seminars: two at UT, one at UL.

We do not provide an analysis of signature courses because we have defined Verification as a grouping for the courses that focus on the wide range of formal techniques that can be used to provide guarantees on the correctness of software systems.

Concerning the most frequently occurring keywords (Table 3, last column), the most prominent ones are ‘program’ and ‘system’, which denote the element that is being verified. These are followed by the generic ‘technique’, and the SE specific ‘software’ against which properties are verified. Other important keywords are ‘security’ and ‘model checking’, showing one of the key goals of verification and a fundamental technique for verifying software and systems.

5.15. PL design

There are 24 courses in the KA PL Design; half at the BSc level, half at the MSc level. The BSc courses are offered by nine universities (except for TU/e), while the twelve MSc-level courses are offered by only six universities, with RU and UU offering four and three courses each. This highlights the different emphasis put on this KA, especially at the MSc level. Only nine courses are taught in SE-specific tracks or programs.

Just like for Verification, we do not provide signature courses. Nonetheless, we observe that six courses focus on compiler construction (**RU**, **RUG**, **TUD**, **UU**, **UvA**, **VU**), four courses address concepts of programming languages (**OU**, **TUD**, **UL**, **UU**), and two courses concern domain-specific languages (**UU**, **TU/e**).

When analysing the most recurring keywords in Fig. 3, after the obvious ‘language’, ‘programming language’, and ‘program’, we find




Fig. 3. Word cloud visualisation of the topics touched upon by the courses that relate to the KA PL Design.




Fig. 4. Correlations between the KAs in the analysed courses, and course clusters. Some KA names are shortened: Req: Requirements; Arch: Architecture; Constr: Construction & Programming; Oper: SE Operations; Mainten: Maintenance; Config: Configuration management; Manag: SE management; Verific: Verification. P.L.D.: Programming language design, C&P: Construction and Programming.

‘type’, which refers to the study of type theory, a fundamental pillar in PL Design (Zaytsev, 2017). This is followed by ‘semantic’, as language semantics is a fundamental topic. The word cloud highlights then ‘compiler’ (PL Design informs compiler construction) and ‘concept’ (but this is originating from the name of some courses).

6. Correlations (RQ₂)

After having analysed the details per KA, we now address RQ_2 by studying the correlations between the KAs. We have determined these by calculating Spearman correlation (we could not assume the data are normally distributed) based on the KAs that were ticked for each course. The raw results are in the online appendix; an overview of the strongest positive correlations is provided in Fig. 4, with the line width and brightness denoting the strength of the correlation between two KAs. Below, we discuss the main clusters of courses based on the existence of stronger correlations between them (conventional interpretations for correlation (Evans, 1996): negligible when $\rho \in [0.00, 0.20]$, weak when $\rho \in [0.20, 0.40]$, moderate when $\rho \in [0.40, 0.60]$, strong when $\rho \in [0.60, 0.80]$, and very strong when $\rho \in [0.80, 1.00]$). Note that, in this context, even weak correlations are indicators of KAs that co-occur often: for a strong correlation to exist, most courses shall cover *both* KA1 and KA2, and this is unlikely, especially in specialised courses that cover a single KA. All the correlation values are shown in the appendix at the end of this paper, in Table 4.

6.1. Requirements, architecture, design, SE models

The KA Requirements has weak correlations with the KAs Design (9/21 Requirements courses also cover Design, $\rho = 0.34$), SE models (9/21, $\rho = 0.30$); and Architecture (7/21, $\rho = 0.26$). This can be explained by the fact that architecture and design are generally seen as phases that follow RE in software development (Nuseibeh, 2001). The co-occurrence of SE models indicates that modelling is an important facet of RE. This is not only testified by the rich literature on requirements modelling, but also by the fact that professional training organisations such as the International Requirements Engineering Board (IREB) do offer certifications on the KA in their advanced level materials (van Akkeren et al., 2016).

Within the seven signature courses, a different situation is portrayed: none of them teaches also Architecture, one of them focuses on Design, and two of them cover SE models. This shows that, while the Requirements KA is often co-taught with other topics in general courses, the more specific Requirements courses are less focused on the relationships with other areas of SE.

In addition to the weak correlation with the above-mentioned KA Requirements, the KA Architecture has a weak correlation with Design (9/22 courses, $\rho = 0.31$) and with SE models (9/22, $\rho = 0.27$). Both terms appear in the second column of Table 2. This is not surprising, as both Design and Architecture focus on design abstractions, with the main difference that Architecture focuses on an additional level that goes beyond the organisation of code and, as explained by Rozanski and Woods (2012), that bridges the gap between the problem and solution spaces.

The KA Design has a moderate correlation with SE models (13/24 courses, $\rho = 0.42$), in addition to the above-mentioned weak correlation with Requirements (9/24, $\rho = 0.34$) and Architecture (9/24, $\rho = 0.31$). The correlation with SE models is unsurprising, as software design is often taught by means of the UML modelling language.

The KA SE models moderate correlation with Design (13/28 courses, $\rho = 0.43$) is witnessed by the existence of books that explain software design via modelling, such as that by Gomaa (2011). The weak correlations with Requirements (9/28, $\rho = 0.30$) and with Architecture (9/28, $\rho = 0.27$) show the typical positioning of modelling as a communication tool that informs the activities before the construction of software systems, in addition to the model-driven SE paradigm (Brambilla et al., 2017).

Finding 7: Requirements, Architecture, Design and SE models are a tightly coupled set of KAs that represent the process for detailing software systems: stakeholders' requirements $\Rightarrow$ system-level architecture $\Rightarrow$ detailed design decisions, often using models to represent those aspects.

6.2. Testing, security, verification

The cluster of the three KAs refers to well-known knowledge in computing education (Garousi et al., 2020): in order to make a system more secure, verification and testing are key activities, with the former concerned with ensuring security in provable manners, and the latter focused on providing probabilistic evidence when full coverage via formal models is not feasible. Despite the three KA of this cluster are linked, we observe that their correlation is overall *more loose than the other clusters* (hence, we do not identify a finding), displaying that, albeit commonalities shared by the three KAs, the topics are often addressed in silos.

The Testing KA (19 courses) has weak correlations with Verification (7 shared courses, $\rho = 0.21$) and a negligible-to-weak one with Software Security (4 shared courses, $\rho = 0.15$). Furthermore, Software Security (16 courses) has a weak correlation with Verification (28 courses, 6 shared with Software Security, $\rho = 0.21$).

6.3. SE operations, SE management, SE process, maintenance, configuration management

The final KA-based cluster of courses concerns different facets of the DevOps paradigm. This cluster results to be only marginally linked to the other clusters, with only few weak correlations between SE management and SE economics ($\rho = 0.21$) and SE Operations and Architecture ($\rho = 0.20$).

Inside the cluster, we observe that the topics are the most highly cohesive across all clusters, with moderate correlations being between SE process and SE Operations ($\rho = 0.41$), SE process and SE management ($\rho = 0.42$), and SE management and SE Operations ($\rho = 0.42$).

Overall, by studying the correlations within this cluster, we can conclude that the topics related to DevOps are seldom taught singularly across all considered universities. We conjecture that this could be either due to (i) the impossibility of treating the theoretical concepts of the KA alone without considering the other KAs, (ii) a general lack of depth with which the topics are treated, leading to the need of considering multiple KA to fill in the content of courses, or (iii) the dissection of a higher concept considered in the course (e.g., DevOps) into finer-grained KA by the SWEBOK.

Finding 8: Because of their interconnectedness, DevOps topics are rarely taught singularly (e.g., via signature courses), but rather covered in courses that touch upon multiple related KAs: SE Operations, SE management, SE process, Maintenance, Configuration Management.

6.4. Lone riders: Construction & programming, PL design, and SE economics

Three KAs appear to be almost completely isolated from the others, namely SE economics, PL Design, and Construction & Programming, although potentially due to different reasons.

In the case of Construction & Programming, the topic treated results to be of a peculiar nature, leading to it need to be considered independently to other KAs. More specifically, a key kind of course in Construction & Programming is the introductory programming course of many SE curricula, which lays the basis on which students can then follow up upon by considering more specialised topics. Therefore, as introductory course, more advanced SE topics are very seldom considered in depth during courses covering this core KA.

Regarding PL Design, instead, its detachment from other KAs might more be due to the theoretical and formal nature of the topic considered, leading it to be only marginally considered in isolation during SE-centric educational curricula.

Finally, the SE economics course is the only that presents a weak correlation to other KA, namely SE management ($\rho = 0.21$). As for the other KA considered above in this section, we conjecture that the peculiar nature of the topic, which focuses more on economic reasoning rather than technical SE aspects, leads the topic to be often treated in specialised courses.

7. University educational foci (RQ₃)

In this section, we address RQ₃ by reporting the results related to the educational focus provided at the ten Dutch universities considered. The distribution of the overall SE education across KAs at the different universities is depicted in Fig. 5. The figure shows that, while several commonalities can be identified across the different universities, some specificities exist too.

In terms of commonalities, most universities result to provide a high number of courses providing a Construction & Programming component, with the UvA providing the highest recurrence of such topic across its SE educational offering. Rather than a specific pedagogical focus on such KA, we deem this high recurrence due to the course *Introduction to programming* provided at the UvA, which is duplicated across

multiple curricula at that university. Overall, the high frequency of the Construction & Programming KA across universities can could be explained by the foundational nature of such KA, especially prevalent at the BSc level in our dataset.

Other commonalities emerge for less recurrent KAs, such as Architecture, Verification, and SE economics. The specialised and advanced nature of such KAs results in their overall low recurrence across all universities, with the exception of UT, that seems to provide a higher rate of specialisation courses on verification topics.

Some KA result instead to be treated with different intensity at the universities considered. For example, the presence of the Software Security KA, while never constituting a considerable portion of SE curricula, varies from a limited fraction of courses (RU) to a notable portion of the educational offerings (TUD). Similarly, a notable portion of courses at RU focus on the PL Design KA, while the topic results to be less recurrent at other institutions (e.g., UT and VU).

Some KA result to be absent at certain institution. Rather than the complete lack of such KAs, we associate such finding to a limited coverage of related topics at the considered institutions, that led to their low representation in the course syllabi, and hence were not included via our data collection process (see Section 4.2).

Overall, we note that all educational contexts provide an extensive and heterogenous range of educational offers, with no institution delivering a narrow and specialised focus on a specific SE topic. The vast range of topics covered in the educational curricula of the considered universities, jointly with the absence of specific pedagogical foci, might be a pointer to the multifaceted skills required to train modern software engineers, which are necessary regardless of specific SE specialisations provided at the universities (e.g., the *Software Engineering and Green IT* MSc track at the VU).

As further note, there are significant global trends that may reshape the SE educational landscape. The popularisation of large language models (LLMs) alone might have a noticeable ripple effect on the educational offerings of universities, from the introduction of new standalone, specialised Bachelor programs in fields such as Generative AI, to the restructuring and redesign of existing SE courses and university curricula. We acknowledge that the educational snapshot of the current state of SE education in the Netherlands might soon be exposed to a some degree of restructuring. Additionally, while our data focuses on computer science programs because SE courses are predominantly embedded there in the Netherlands, it would be interesting to explore degrees beyond “classical” computer science, such as in AI or data science, and compare the results.

Finding 9: While minor differences exist for specific KAs (e.g., Software Security and PL Design), the analysed universities provide a high degree of commonality in terms of the covered KAs, indicating that the studied KAs are necessary traits of modern software engineers. We expect more noticeable differences in terms of the specialisation focus, e.g., sustainability or artificial intelligence. These aspects are, however, not covered by the SWEBOK, the theoretical lens of our analysis. The rapid adoption of LLMs might soon have a noticeable ripple effect on SE education, leading to the restructuring and redesign of existing courses and university curricula.

8. Related work

Lee and Cheng (2011) showed that Taiwan faces significant challenges in SE education, with issues in both the quantity and quality of prepared graduates. Despite a surplus of IT-related students, many lack core competencies. This is attributed to an overemphasis on introductory courses and neglect of advanced topics like software quality, modelling, and domain knowledge. The Consortium of Taiwan has addressed these gaps by introducing a module-oriented curriculum




Fig. 5. Heatmap across KAs and universities: percentage of courses covering a specific KA within a given university.

inspired by SWEBOK, leading to increased course offerings and enrolment. However, persistent dysfunctions, such as avoidance of process and insufficient attention to practical domain applications, still hinder the alignment of education with industry needs. Our research extends these efforts by analysing SE education in the Netherlands, focusing on the breadth and co-occurrence of KAs across graduate and undergraduate courses. Unlike previous studies, we analyse interdependencies among KAs and assess variations in educational focus across universities. This approach offers a comparative perspective to identify unique strengths and gaps in preparing graduates for SE challenges. Despite the lack of ambition to country-wide curriculum revision, we offer comparable insights gained from comparing existing SE courses in the Netherlands, and aligning them with SWEBOK chapters.

In their systematic review, Tenhunen et al. (2023) analysed 127 papers related to SE capstone courses (2007–2022). They found that capstone courses in computer science and SE programs aim to prepare students for professional life by combining technical skills with essential soft skills such as teamwork, communication, and project management. They identified common features of SE capstone courses, such as semester-long duration, team sizes of 4–5 students, and the use of real-world projects, often with external clients. While these courses effectively cover software development life-cycle stages, gaps remain in areas like software maintenance and the integration of continuous peer assessment. While this systematic review provides a comprehensive understanding of current SE capstone courses, our research sets itself apart by focusing specifically on the educational landscape within the Netherlands.

Huang et al. (2021) presented a tertiary study that synthesised and classified 26 SLRs published between 2004 and 2019. They noted that while several individual SLRs have focused on specific aspects of SE education, none have provided a broad, high-level overview of the entire SE education research landscape. By offering insights into the frequently applied methods and tools, as well as the recurring themes and issues in SE education, they provided a broader and more holistic view of the field. Additionally, they identified key challenges in SE

education that have emerged from these studies which can help to inform future research and development in SE pedagogy.

Malik and Zafar (2012) conducted a mapping study focusing on improving SE education practices. This study provided an overview of the literature addressing SE education advancements, identifying and categorising 70 primary studies published according to a well-defined SE educational framework. The authors highlighted that inadequate curricula for SE remain a critical risk, emphasising the need for curricula that align with industry demands and equip students with the necessary skills, knowledge, and expertise. The study mapped primary studies to reference curricula such as SWEBOK, identifying gaps in these frameworks that could benefit from revision to better reflect diverse aspects of SE. Additionally, they analysed publication trends, noting that conferences like CSEE&T and regions like the USA and China lead in SE educational research. The authors proposed that their findings can inform curriculum refinement, align academic practices with industry needs for practitioners, and inspire new research directions for researchers.

Qadir and Usman (2011) presented a systematic mapping study to synthesise and aggregate literature on SE curriculum design, revision, and assessment, providing a comprehensive overview of efforts reported in conferences like CSEE&T, FIE, and EDUCON. Their study highlighted the historical evolution of SE as a discipline, emphasising the transition from its early conception as a subfield of Computer Science to a mature academic field with stand alone programs at BSc and MSc levels. The authors noted that contributions related to course-level improvements far outnumber those focused on programme-level enhancements, signalling a need for greater emphasis on holistic curriculum development. They underlined key frameworks such as SWEBOK, SE 2004, and GSwE2009, which have guided SE curriculum advancements, and also identified a growing need to align curricula with industry trends and demands. Our research builds on this work by exploring the state of SE curricula within the Netherlands, focusing on the co-occurrence and integration of knowledge areas across various academic institutions. Unlike this mapping study, which primarily categorises contributions by quantity and type, our approach evaluates the interdependencies of curriculum components and their practical alignment with industry requirements, offering actionable insights into the broader landscape of SE education.

Liargkovas et al. (2021) investigated the gap between software practitioners' education, as outlined in the IEEE/ACM SE2014 SEEK, and industrial needs, using Wikipedia articles cited in Stack Overflow posts as a proxy for developers' real-world informational needs. Their study revealed that while SEEK adequately covers foundational topics like computer science fundamentals, software design, and mathematical concepts, it lacks sufficient emphasis on practical areas such as the World Wide Web, testing, security, and soft skills. The authors identified significant misalignments between the SEEK curriculum and industry requirements, suggesting that future curriculum models should integrate industry-driven topics (e.g., software testing, design, and security) and soft skills (e.g., teamwork, collaboration, and project management) to better prepare students for professional challenges. Our research complements these findings by focusing on curriculum design within the Netherlands, evaluating the integration and co-occurrence of knowledge areas across universities to identify unique gaps and opportunities for alignment with industrial demands.

Some studies analysed the academia-industry gap through surveys with SE graduates. Lethbridge (2002)'s seminal work was the first one to explore the gap in terms of which topics were considered important, were learned in education, and/or learned on the job. Lethbridge's instrument was later adapted by Kitchenham et al. (2005), who conducted a similar survey with thirty graduates of four universities in the United Kingdom. Our study focuses on analysing the course descriptions, rather than surveying graduates, although we reckon that triangulating with graduates' perception would be interesting.

A systematic mapping study on Software Engineering for Sustainability (SE4S) provided an updated and in-depth overview of the field, investigating recent contributions, knowledge areas, and research facets (Penzenstadler et al., 2014). They offer a research map that identified trends, tools, methods, and frameworks. The findings revealed that while SE4S has grown significantly, its impact remains concentrated in a few knowledge areas, highlighting gaps that require further exploration. The authors emphasised the need for a roadmap to guide future SE4S research and development. This study aligns with broader efforts to adapt SE practices to address sustainability challenges, offering insights for aligning curricula and research agendas with emerging societal needs.

In contrast to prior research, including studies focused more narrowly on evidence-based SE practices (Kitchenham et al., 2007; Liu et al., 2021), and those offering a broader, more holistic view of SE education (Huang et al., 2021; Malik and Zafar, 2012), our study serves as a reference point for academics and practitioners looking to understand SE education in the Netherlands. By examining the co-occurrence and integration of KAs across undergraduate and graduate programs, we provide a nuanced perspective on how Dutch institutions align their curricula with industry requirements and emerging trends in SE. This research not only highlights gaps and opportunities within the Netherlands' educational landscape but also offers actionable insights into curriculum design and interdependencies among KAs, bridging the divide between academic preparation and professional practice.

9. Discussion

We analyse the findings of our study in relation to the research questions, highlighting key insights and their significance. We begin by addressing the research questions and interpreting the results in the context of our study. Next, we discuss potential threats to validity and limitations, acknowledging factors that may have influenced our outcomes. Finally, we explore the broader implications of our findings for SE education, providing insights that may inform future research and practice.

9.1. Answers to the research questions

RQ₁: What are the topics studied in SE Courses?

We found that the various KAs are generally covered to a reasonably comparable extent (see Fig. 2). While the KA Construction & Programming is the most frequent (74 courses, while the second most common, SE models, has 28 courses), this can easily be explained by the fact that programming is a basic skill for both computer science and SE curricula, also demonstrated by the fact that 62/74 courses in the Construction & Programming KA are at the BSc level. Some differences can be seen when comparing KAs across BSc and MSc: for example, Maintenance is mostly covered by MSc courses (18/19), while other KAs are perfectly balanced, like SE models (14 per level), Design (12 per level), and PL Design (12 per level). The presence of a high number of courses covering the KA Verification is a remarkable finding, which is in line with the research tradition in the Netherlands. A couple KAs have only nine courses (less than one per university); these are Configuration Management and SE economics. The word clouds included in Section 5 provide a more in-depth analysis for the reader interested in learning the most frequent topics mentioned in the course descriptions.

RQ₂: What are the (co-)occurrences among studied topics and the educational context?

We conducted a correlation analysis of the KAs covered by the courses (see Fig. 4). The results highlight three clusters of KAs that are tightly coupled (i.e., courses often cover more than one of such KAs), explained in Sections 6.1–6.3: (i) SE activities that pertain to the requirements analysis and design (Requirements, Architecture, Design, SE models); (ii) techniques and tools for quality assurance,

including security (Testing, Software Security, Verification); and (iii) process-oriented and DevOps-related topics (SE Operations, SE management, SE process, Maintenance, Configuration Management). The analysis also revealed some isolated KAs that are generally taught on their own: these are Construction & Programming, PL Design, and SE economics. The former two can be explained by the fact these topics are part of general computer science programs; for SE economics, we speculate this occurs due to the existence of very specialised courses that focus on the business and managerial aspects of SE.

RQ₃: Do universities differ in terms of education foci? Our analysis of the number of courses covering each KA at the various institutions (see Fig. 5) does not reveal significant discrepancies: Dutch universities seem to cover all the analysed KAs via their courses. The most remarkable outlier is the high number of courses on Construction & Programming at UvA, but this is explained by the existence of numerous variants of programming courses that are offered to various programs across the university. Smaller deviations exist, but they simply reflect specificities of some universities, often originating from their research foci.

9.2. Threats to validity and limitations

In this section, we discuss the most relevant threats to validity that might have influenced our research, by considering common pitfalls and suggestions to consider threats in empirical studies (Lago et al., 2024). Additionally, we document the limitations our research results entail.

9.2.1. Internal validity

A potential internal threat to our results is constituted by the labelling strategy utilised to code the courses according to the predefined KA. To mitigate this threat, we relied on a crowdsourcing strategy by involving educators of all universities to tag their own courses (see Phase 4 of our research method in Section 4.2). The final tags resulting from the homogenisation strategy (Phase 5) were then revised again by the original annotators, in order to further mitigate internal threats arising from the homogenisation process. The tagging reached *substantial* agreement $\kappa = 0.718$, although we acknowledge that a higher degree of agreement would be desirable.

In order to mitigate potential threats related to the data analysis, only simple summary statistics are utilised to present the results, while reporting as much as possible the raw data on which our results are based (e.g., the number of courses considering a specific KA).

9.2.2. Construct validity

An inherent threat to the validity of our results lies in the predefined Knowledge Areas (KAs) used for data collection. Specifically, the KAs outlined in the SWEBOK guide might not fully capture the nuanced topics of the analysed software engineering university courses. Relying on an *a priori* defined catalog of KAs was deemed however a necessary tradeoff between construct validity and internal validity. While such research design choice potentially sacrificed some expressiveness and simplified the studied construct, it helped mitigate subjective bias during the data analysis process. As overarching mitigation strategy of this threat, we designed a tailored research process to refine and adapt the SWEBOK KAs specifically for our study's context (see Phase 3 of the research method in Section 4.2).

Regarding a further simplification of the observed construct that might have influenced our results, in our study we do not go into the details on study program structures, e.g., by considering mandatory versus optional courses, study credits associated to courses, prerequisite course ordering, nature of the thesis project, assigned grades, and student satisfaction questionnaires. While considering these characteristics would have yielded to a more nuanced understanding of the educational landscape, doing so for all 207 courses would have

been prohibitively challenging, and opening to threats such as data deluge and analysis paralysis. The design choice of not considering such information was hence a necessary tradeoff between construct validity and both internal validity and study feasibility. This limitation can be addressed in future work by designing a study specifically tailored to focus on these more granular facets.

Finally, another potential threat to construct validity regards the original source of data of our study, namely the syllabi of the courses considered. In fact, the content of the syllabi, which are usually drafted before the start of courses, might deviate from the actual content of the live lecture. Using the syllabi as original data source entailed a trade-off between construct and external validity, which allows us to consider a rather high number of courses, namely 207, that could not be considered if another type or research method, e.g., interviews with educators, would be adopted. To mitigate this threat, researchers label exclusively syllabi of their own university, resorting to contacting the educator in charge of the course whenever the information on the syllabi is deemed insufficient.

9.2.3. External validity

The results presented in this study are by definition bound to the specific geo-educational context considered, namely SE higher education in the Netherlands. While the results and conclusions might be applicable to different extents to other countries, we cannot not claim the generalisability of the results. While the implications presented in Section 9.3 might be corroborated by considering similar pedagogical contexts to the one chosen for this study, further research is needed to prove such assertion.

At the intersection of external and construct validity, we warn interested readers and researchers that the research process, keywords, and resulting KA might be applicable and replicable only in similar educational and societal contexts, namely SE university courses with frontal lectures in Western countries. Further research should be conducted to study if, and in affirmative case to what extent, the research process and intermediate research artefacts can be transposed to different pedagogical contexts and paradigms such as peer learning.

9.2.4. Limitations

As limitations to our work, due to the volume of courses considered, we are not able to study in depth many nuanced facets of the courses considered, such as topics covered within each KA, student grades, and number of educators involved in each course among others. Therefore, the results and conclusions presented in this study need to be interpreted exclusively as a higher level overview of the educational offerings of SE universities in the Netherlands. Numerous related concepts, such as student perceptions, impediments encountered by educators, graduate employability, industry expectations, and effectiveness of different pedagogical components, and variation of the educational landscape in time, remain uncovered with this research and are left up for future work.

9.3. Implications for SE education

The findings suggest actionable insights that can be used to enhance and diversify SE curricula.

Integrative courses on software design activities. The pairwise correlations between the KAs Requirements, Design, Architecture, and SE models indicate that there may be an opportunity to foster a more integrated teaching approach across these KAs. Educators could consider designing broader courses that allow students to engage with these KAs holistically, with an emphasis on how they inform each other throughout the software development lifecycle. This could be done via the creation of modules that consist of sub-modules, each focusing on one specific KA, with an integrative project where the students would learn about the complex dependencies across the KAs via exercising their learned knowledge (and techniques) on a specific case.

Software security via testing and verification. Similarly, the correlations among Testing, Software Security, and Verification provide another opportunity for educators to create integrated learning experiences. Given the increasing importance of security and reliability in software systems, educators could develop projects or case studies that require students to apply verification methods and testing techniques to ensure security in software systems. This could be particularly effective in advanced MSc courses, where students can explore the theoretical underpinnings of these KAs and their application in practice. Additionally, offering seminars or project-based courses focused on security and verification could help students gain hands-on experience with state-of-the-art tools and methodologies.

Bridging the SE economics island. The KA SE economics is not only under-represented in the analysed SE curricula, but it is also isolated from the rest. As this area is highly relevant to the industry, and for students to comprehend the business aspects of software, educators should consider expanding the scope of SE programs by integrating courses that address the economic aspects of software development, including business models, software ecosystems, and sustainability. Our analysis reveals that these topics, when taught, are more often embedded at the MSc level; educators could consider introducing them earlier in the curriculum to better prepare students for the job market. Interdisciplinary courses combining SE with business, economics, and entrepreneurship could be particularly valuable in this regard.

10. Summary and future work

This paper provided a systematic analysis of the SE higher education landscape in the Netherlands, through an analysis of 207 courses offered by ten universities. We have identified that (RQ_1) the KA Construction & Programming is unavoidably the most frequent, followed by a quite balanced distribution across the other KAs, with the exception of Configuration Management and SE economics. Our analysis of co-occurrences (RQ_2) revealed clusters of KAs that are often co-taught, regarding analysis and design, quality assurance, and SE process and DevOps. Finally, we could not find significant differences in the distribution of the KA coverage across universities (RQ_3).

Our findings point to a need for further exploration into the teaching of SE economics, a less represented KA that appears to be highly relevant to industry but still under-emphasised in academic settings. Researchers could investigate how to better integrate economic principles such as software production models, business ecosystems, and sustainability into SE curricula, especially at the BSc level. This could involve developing interdisciplinary courses or modules that combine SE with business and entrepreneurship.

The findings also suggest several important areas for future research, particularly in the context of the correlations between different KAs and their integration into SE education. One significant observation is the strong correlation between the KAs of Requirements, Design, Architecture, and SE models. This highlights the necessity for further research on how these KAs interact in both academic curricula and industry practices, with a focus on understanding the dependencies and synergies between them. Researchers should explore how these correlations can be leveraged to create more integrated learning modules or courses that bridge gaps between foundational concepts like Requirements and advanced topics like Architecture and Design. The strong ties between Requirements and Design, as well as the role of SE models as communication tools in the development lifecycle, suggest that more research is needed to understand how these KAs can be effectively taught together, potentially through model-driven SE paradigms.

Moreover, the cluster involving Testing, Software Security, and Verification offers an opportunity to explore the interconnected nature of these KAs in the context of ensuring software quality and security. Researchers could investigate how the principles of verification and testing can be aligned with security practices in educational settings,

particularly through cross-KA projects that allow students to apply formal verification methods, testing strategies, and security principles simultaneously. This would not only enrich students' practical skills but also enhance their understanding of the complexities involved in developing secure and reliable software systems. Research on the alignment between Verification and PL Design is also promising, as both fields share formal mathematical foundations and could benefit from joint curricula focusing on formal methods, type systems, and language semantics.

A useful direction would be extending the study by considering the ordering of the courses and possible pre-requisites, in an attempt to study the progression across KAs within specific study paths, besides the correlations. Furthermore, there are important trends that are affecting computing education: while the SWEBOK is a robust resource for the mapping, it is rather conservative and does not cover emerging areas like Generative AI (Daun and Brings, 2023), Quantum computing (Murillo et al., 2025), sustainability (Moreira et al., 2025), etc. Future studies could take these topics into account as well.

As future work, we plan to replicate our nationwide study within the Italian software engineering higher education system, where one of the authors is currently based. Conducting a study of this scale demands several prerequisites, including established institutional partnerships, deep familiarity with the local educational system, and the cooperation of at least one educator from each participating university. While meeting all these criteria simultaneously is challenging, we highly encourage researchers to repeat our study in their national context to uncover regional similarities and differences. Furthermore, we invite international comparative studies, though scholars have eloquently noted the substantial challenges such cross-border efforts entail (Teichler, 2014).

CRedit authorship contribution statement

Bastiaan Heeren: Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Fabiano Delpiaz:** Writing – review & editing, Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Mazyar Seraj:** Writing – review & editing, Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Roberto Verdecchia:** Writing – review & editing, Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Vadim Zaytsev:** Writing – review & editing, Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We dedicate this work in loving memory of Bastiaan Heeren, who passed away while this work was being completed. His passion and dedication to software engineering education will not be forgotten, and we are proud and honoured to have had the opportunity to collaborate with him. We would also like to thank all those who contributed to and helped with the collection and verification of the data, especially Marcello Bonsangue (University of Leiden), Andrea Capiluppi (University of Groningen), Patricia Lago (Vrije Universiteit Amsterdam), Ana Oprescu (University of Amsterdam), Annibale Panichella (Delft University of Technology).

Table 4

Correlations between the KAs (via Spearman). The colour coding is based on the standard interpretation by Evans (1996): very strong ($|r_s| \geq 0.8$): **Dark Green**; strong ($0.6 \leq |r_s| < 0.8$): **Medium Green**; moderate ($0.4 \leq |r_s| < 0.6$): **Light Green**; weak ($0.2 \leq |r_s| < 0.4$): **Pale Green**; negligible ($|r_s| < 0.2$): **Pale Yellow**. For negatives, similar shading but in tones of red.

	Requirements	Architecture	Design	Construction & Programming	Testing	SE Operations	Maintenance	Configuration Management	SE management	SE process	SE models	Software Security	SE economics	Verification	PL Design
Requirements	1.00	0.26	0.34	-0.11	0.01	0.03	-0.05	-0.07	0.09	0.13	0.30	-0.09	-0.07	-0.08	-0.12
Architecture	0.26	1.00	0.32	-0.19	-0.11	0.21	0.11	0.08	0.12	-0.08	0.28	-0.10	0.08	-0.14	-0.12
Design	0.34	0.32	1.00	0.01	-0.01	0.07	0.04	0.14	0.06	0.14	0.43	-0.10	-0.08	-0.14	-0.13
Construction & Programming	-0.11	-0.19	0.01	1.00	-0.14	-0.17	-0.20	-0.01	-0.24	-0.13	-0.03	-0.03	-0.11	-0.21	-0.08
Testing	0.01	-0.11	-0.01	-0.14	1.00	0.04	-0.04	0.01	-0.12	0.04	0.17	0.16	-0.07	0.22	-0.12
SE Operations	0.03	0.21	0.07	-0.17	0.04	1.00	0.23	0.21	0.42	0.41	0.05	-0.08	0.03	-0.11	-0.10
Maintenance	-0.05	0.11	0.04	-0.20	-0.04	0.23	1.00	0.26	0.30	0.34	-0.03	-0.03	0.10	-0.13	-0.12
Configuration Management	-0.07	0.08	0.14	-0.01	0.01	0.21	0.26	1.00	0.07	0.36	-0.02	-0.06	-0.05	-0.08	-0.08
SE management	0.09	0.12	0.06	-0.24	-0.12	0.42	0.30	0.07	1.00	0.42	-0.05	-0.10	0.22	-0.14	-0.13
SE process	0.13	-0.08	0.14	-0.13	0.04	0.41	0.34	0.36	0.42	1.00	0.07	-0.05	0.14	-0.10	-0.13
SE models	0.30	0.28	0.43	-0.03	0.17	0.05	-0.03	-0.02	-0.05	0.07	1.00	-0.06	0.05	-0.03	-0.10
Software Security	-0.09	-0.10	-0.10	-0.03	0.16	-0.08	-0.03	-0.06	-0.10	-0.05	-0.06	1.00	-0.06	0.20	0.01
SE economics	-0.07	0.08	-0.08	-0.11	-0.07	0.03	0.10	-0.05	0.22	0.14	0.05	-0.06	1.00	-0.08	-0.08
Verification	-0.08	-0.14	-0.14	-0.21	0.22	-0.11	-0.13	-0.08	-0.14	-0.10	-0.03	0.20	-0.08	1.00	0.08
PL Design	-0.12	-0.12	-0.13	-0.08	-0.12	-0.10	-0.12	-0.08	-0.13	-0.13	-0.10	0.01	-0.08	0.08	1.00

Appendix. Correlations between the KAs

See Table 4.

Data availability

All data used for this research, including the coded data and the scripts used for our analysis is available online as a companion replication package of our work, available at: <https://doi.org/10.5281/zenodo.21276768>.

References

- Akdur, D., 2022. Analysis of software engineering skills gap in the industry. *ACM Trans. Comput. Educ.* 23 (1), <http://dx.doi.org/10.1145/3567837>.
- van Akkeren, E., Baumann, L., Cannegieter, J.J., Hood, C., Hruschka, P., Lampe, M., Leutbecher, E., van Loenhoud, H., de Roo, P., Staal, S., et al., 2016. Handbook of requirements modeling according to the IREB standard. IREB Int. Req. Eng. Board.
- Ardis, M., Budgen, D., Hislop, G.W., Offutt, J., Sebern, M., Visser, W., 2015. SE 2014: Curriculum guidelines for undergraduate degree programs in software engineering. *Computer* 48 (11), 106–109.
- Basili, V., Caldiera, G., Rombach, H.D., 1994. The goal question metric approach. *Encycl. Softw. Eng.* 528–532.
- Boehm, B.W., Sullivan, K.J., 2000. Software economics: A roadmap. In: *Proceedings of the Conference on the Future of Software Engineering*, pp. 319–343.
- Borman, K.M., Clarke, C., Cotner, B., Lee, R., 2012. Cross-case analysis. In: *Handbook of Complementary Methods in Education Research*. Routledge, pp. 123–139.
- Brambilla, M., Cabot, J., Wimmer, M., 2017. Model-Driven Software Engineering in Practice, second ed. Morgan & Claypool Publishers, <http://dx.doi.org/10.1007/978-3-031-02549-5>.
- van den Broek, A., van Casteren, W., 2018. Tracks in het hoger onderwijs. ResearchNed, http://www.researchned.nl/wp-content/uploads/2019/01/Tracks-in-het-hoger-onderwijs_def.pdf.
- Burgueño, L., Ciccozzi, F., Famelis, M., Kappel, G., Lambers, L., Mosser, S., Paige, R.F., Pierantonio, A., Rensink, A., Salay, R., Taentzer, G., Vallecillo, A., Wimmer, M., 2019. Contents for a model-based software engineering body of knowledge. *Softw. Syst. Model.* 18 (6), 3193–3205.
- Chitchyan, R., Mahmoudi, N., 2025. Embedding sustainability in software engineering curriculum: A case study. URL <https://arxiv.org/abs/2510.03321>, arXiv:2510.03321.
- on Computing Curricula, T.J.T.F., 2024. Computer Science Curricula 2023: The Final Report. Technical Report, ACM, IEEE-CS, AAAI, <https://csed.acm.org/cs2023-report-with-feedback/>.

- Cornelisse, S., Boor, I., et al., 2022. De kracht van samenwerkende docententeams bij curriculumafstemming. <https://tlc.uva.nl/article/zichtbare-leerlijnen-programma/>.
- Cusumano, M.A., 2008. The changing software business: Moving from products to services. *Computer* 41 (1), 20–27. <http://dx.doi.org/10.1109/MC.2008.29>.
- Daun, M., Brings, J., 2023. How ChatGPT will change software engineering education. In: *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*, pp. 110–116.
- Evans, J.D., 1996. *Straightforward Statistics for the Behavioral Sciences*. Thomson Brooks/Cole Publishing Co.
- Galster, M., Angelov, S., 2016. What makes teaching software architecture difficult? In: *Proceedings of the 38th International Conference on Software Engineering Companion*, pp. 356–359.
- Garousi, V., Rainer, A., Lauvås, Jr., P., Arcuri, A., 2020. Software-testing education: A systematic literature mapping. *J. Syst. Softw.* 165, 110570.
- Gary, K., 2015. Project-based learning. *Computer* 48 (9), 98–100. <http://dx.doi.org/10.1109/MC.2015.268>.
- Gilgun, J.F., 2019. Deductive qualitative analysis and grounded theory: Sensitizing concepts and hypothesis-testing. *SAGE Handb. Curr. Dev. Grounded Theory* 55, 107–122.
- Glaser, B.G., Strauss, A.L., et al., 1998. *Grounded theory. Strat. Qual. Forsch. Bern*: Huber 4.
- Gomaa, H., 2011. *Software Modeling and Design: UML, Use Cases, Patterns, and Software Architectures*. Cambridge University Press.
- Hassan, A.E., Oliva, G.A., Lin, D., Chen, B., Jiang, Z.M., 2024. Towards AI-native software engineering (SE 3.0): A vision and a challenge roadmap. *ACM Trans. Softw. Eng. Methodol.*
- Huang, X., Zhang, H., Zhou, X., Shao, D., Jaccheri, L., 2021. A research landscape of software engineering education. In: *2021 28th Asia-Pacific Software Engineering Conference. APSEC, IEEE*, pp. 181–191.
- Keuzegids, 2024a. Alle masters — Sector exact & informatica — Computerscience. <https://www.keuzegids.org/ol/gidsen/ma24/3913>.
- Keuzegids, 2024b. Alle wo-opleidingen — Sector exact en informatica — Wo-ICT. <https://www.keuzegids.org/ol/gidsen/uni24/3737>.
- Kitchenham, B., Budgen, D., Brereton, P., Turner, M., Charters, S., Linkman, S., 2007. Large-scale software engineering questions—expert opinion or empirical evidence? *IET Softw.* 1 (5), 161–171.
- Kitchenham, B., Budgen, D., Brereton, P., Woodall, P., 2005. An investigation of software engineering curricula. *J. Syst. Softw.* 74 (3), 325–335.
- Kosmützky, A., 2016. The precision and rigor of international comparative studies in higher education. In: *Theory and Method in Higher Education Research*. Emerald Group Publishing Limited, pp. 199–221.
- Kosmützky, A., Krücken, G., 2014. Growth or steady state? A bibliometric focus on international comparative higher education research. *High. Educ.* 67 (4), 457–472.
- Kosmützky, A., Nokkala, T., Diogo, S., 2020. Between context and comparability: Exploring new solutions for a familiar methodological challenge in qualitative comparative research. *High. Educ. Q.* 74 (2), 176–192.

- Lago, P., Runeson, P., Song, Q., Verdecchia, R., 2024. Threats to validity in software engineering—hypocritical paper section or essential analysis? In: Proceedings of the 18th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. pp. 314–324.
- Lee, J., Cheng, Y.C., 2011. Change the face of software engineering education: A field report from Taiwan. *Inf. Softw. Technol.* 53 (1), 51–57.
- Lethbridge, T.C., 2002. What knowledge is important to a software professional? *Computer* 33 (5), 44–50.
- Liargkovas, G., Papadopolou, A., Kotti, Z., Spinellis, D., 2021. Software engineering education knowledge versus industrial needs. *IEEE Trans. Educ.* 65 (3), 419–427.
- Liu, C., Gao, C., Xia, X., Lo, D., Grundy, J., Yang, X., 2021. On the reproducibility and replicability of deep learning in software engineering. *ACM Trans. Softw. Eng. Methodol.* (TOSEM) 31, 1–46. <http://dx.doi.org/10.1145/3477535>.
- Malik, B., Zafar, S., 2012. A systematic mapping study on software engineering education. *Int. J. Educ. Pedagog. Sci.* 6 (11), 3343–3353.
- Manikas, K., Hansen, K.M., 2013. Software ecosystems — A systematic literature review. *J. Syst. Softw.* 86 (5), 1294–1306. <http://dx.doi.org/10.1016/j.jss.2012.12.026>.
- Meijers, A.W.M., Borghuis, V.A.J., Mutsaers, E.J.P.J., Overveld, van, C.W.A.M., Perrenet, J.C., 2005. Criteria for Academic Bachelor's and Master's Curricula, second ed. Technical University of Eindhoven.
- Moreira, A., Lago, P., Heldal, R., Betz, S., Brooks, I., Capilla, R., Coroamă, V.C., Duboc, L., Fernandes, J.P., Leifler, O., et al., 2025. A roadmap for integrating sustainability into software engineering education. *ACM Trans. Softw. Eng. Methodol.* 34 (5), 1–27.
- Murillo, J.M., Garcia-Alonso, J., Moguel, E., Barzen, J., Leymann, F., Ali, S., Yue, T., Arcaini, P., Pérez-Castillo, R., García-Rodríguez de Guzmán, I., et al., 2025. Quantum software engineering: Roadmap and challenges ahead. *ACM Trans. Softw. Eng. Methodol.* 34 (5), 1–48.
- Murphy, E., Crick, T., Davenport, J.H., 2016. An analysis of introductory programming courses at UK universities. *arXiv preprint arXiv:1609.06622*.
- Naik, K., Tripathy, P., 2011. *Software Testing and Quality Assurance: Theory and Practice*. John Wiley & Sons.
- Nuseibeh, B., 2001. Weaving together requirements and architectures. *Computer* 34 (3), 115–119.
- Overheid, 2024. Regeling aanmelding en toelating hoger onderwijs. <https://wetten.overheid.nl/BWBR0035059/2024-09-01>.
- Overheid, 2025. Wet op het hoger onderwijs en wetenschappelijk onderzoek. <https://wetten.overheid.nl/BWBR0005682/2025-01-01/>.
- Penzenstadler, B., Raturi, A., Richardson, D., Calero, C., Femmer, H., Franch, X., 2014. Systematic mapping study on software engineering for sustainability (SE4s). In: Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering. pp. 1–14.
- Perry, D.E., Wolf, A.L., 1992. Foundations for the study of software architecture. *ACM SIGSOFT Softw. Eng. Notes* 17 (4), 40–52.
- Pressman, R.S., Maxim, B.R., 2019. *Software Engineering: A Practitioner's Approach*, ninth ed. McGraw-Hill Education, New York, NY.
- Qadir, M.M., Usman, M., 2011. Software engineering curriculum: A systematic mapping study. In: 2011 Malaysian Conference in Software Engineering. IEEE, pp. 269–274.
- Rozanski, N., Woods, E., 2012. *Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives*. Addison-Wesley.
- Simon, Mason, R., Crick, T., Davenport, J.H., Murphy, E., 2018. Language choice in introductory programming courses at australasian and UK universities. In: Proceedings of the 49th ACM Technical Symposium on Computer Science Education. SIGCSE '18, Association for Computing Machinery, New York, NY, USA, pp. 852–857. <http://dx.doi.org/10.1145/3159450.3159547>.
- Sommerville, I., 2015. *Software Engineering*, tenth ed. Pearson, Boston, MA.
- Soosai Raj, A.G., Ketsuriyonk, K., Patel, J.M., Halverson, R., 2018. Does native language play a role in learning a programming language? In: Proceedings of the 49th ACM Technical Symposium on Computer Science Education. SIGCSE '18, Association for Computing Machinery, New York, NY, USA, pp. 417–422. <http://dx.doi.org/10.1145/3159450.3159531>.
- Studiekeuze, L.C., 2025. Hogeschool of universiteit: wat zijn de verschillen?.
- Subramaniam, V., van der Hagen, T., Smits, R., Heimovaara, S., 2025. 4TU.Federation. <https://www.4tu.nl/>.
- Teichler, U., 2014. Opportunities and problems of comparative higher education research: The daily life of research. *High. Educ.* 67 (4), 393–408.
- Tenhunen, S., Männistö, T., Luukkainen, M., Ihanola, P., 2023. A systematic literature review of capstone courses in software engineering. *Inf. Softw. Technol.* 107191.
- Tramontana, P., Marin, B., Paiva, A.C., Mendes, A., Vos, T.E., Amalfitano, D., Cammaerts, F., Snoeck, M., Fasolino, A.R., 2024. State of the practice in software testing teaching in four European countries. In: 2024 IEEE Conference on Software Testing, Verification and Validation. ICST, IEEE, pp. 59–69.
- University of Twente, 2023. Learning-by-interacting: The university of twente vision on learning and teaching. <https://www.utwente.nl/en/service-portal/organisation-regulations-and-codes-of-conduct/vision-on-learning-and-teaching>.
- Visscher-Voerman, I., Muller, A., 2017. Curriculum development in engineering education: Evaluation and results of the twente education model (TOM). In: 45th SEFI Annual Conference. URL https://ris.utwente.nl/ws/portalfiles/portal/19806823/sefi_curriculum_development_in_engineering_education.pdf.
- Washizaki, H., et al., 2024. Guide to the software engineering body of knowledge, version 4.0. <http://www.swebok.org>.
- Wijngaards, L., Bok, H., Flesch, F., Leplaa, H., Loopers, J., Merx, S., Veenhoven, J., 2020. Zichtbare leerlijnen. <https://www.uu.nl/onderwijs/leerlijnen>.
- Xu, L., Brinkkemper, S., 2007. Concepts of product software. *Eur. J. Inf. Syst.* 16, 531–541. <http://dx.doi.org/10.1057/palgrave.ejis.3000703>.
- Zaytsev, V., 2017. Language design with intent. In: Batory, D., Gray, J., Kulkarni, V. (Eds.), Proceedings of the ACM/IEEE 20th International Conference on Model Driven Engineering Languages and Systems. MODELS, IEEE, pp. 45–52. <http://dx.doi.org/10.1109/MODELS.2017.16>.